

Avoiding bias in reconstructing the largest observable scales from partial-sky data

Stephen M. Feeney,^{1,*} Hiranya V. Peiris,^{1,†} and Andrew Pontzen^{2,‡}

¹*Department of Physics and Astronomy, University College London, London WC1E 6BT, U.K.*

²*Institute of Astronomy and Kavli Institute for Cosmology,
University of Cambridge, Cambridge CB3 0HA, U.K.*

(Dated: September 28, 2011)

Obscuration due to Galactic emission complicates the extraction of information from cosmological surveys, and requires some combination of the (typically imperfect) modeling and subtraction of foregrounds, or the removal of part of the sky. This particularly affects the extraction of information from the largest observable scales. Maximum-likelihood estimators for reconstructing the full-sky spherical harmonic coefficients from partial-sky maps have recently been shown to be susceptible to contamination from within the sky cut, arising due to the necessity to band-limit the data by smoothing prior to reconstruction. Using the WMAP 7-year data, we investigate modified implementations of such estimators which are robust to the leakage of contaminants from within masked regions. We provide a measure, based on the expected amplitude of residual foregrounds, for selecting the most appropriate estimator for the task at hand. We explain why the related quadratic maximum-likelihood estimator of the angular power spectrum does not suffer from smoothing-induced bias.

INTRODUCTION

It is unavoidable that we observe the Universe through the galaxy we inhabit. The foreground contamination injected by the Milky Way into full-sky cosmological datasets must be modeled and removed, or the regions most conspicuously contaminated must be excised. Where no precise model of the foregrounds is available, cutting the sky is the most robust option, with the regrettable consequence that part of the signal is discarded along with the contamination. This includes information on the largest scales, which are valuable for a variety of reasons, including measurement of the integrated Sachs-Wolfe effect [1] and constraining primordial non-Gaussianity using tracers of large-scale structure [2].

It is impossible to uniquely recover the cosmological signal discarded in the sky cut. However, by writing down the likelihood for the region of the sky in which one trusts the data, it is possible to reconstruct an estimate of the signal at large scales which maximizes the likelihood of the residual noise [3]. An alternative reconstruction scheme maximizes the *posterior* probability [4–8] of measuring the underlying cosmological signal given the available data and a prior theoretical expectation on the signal.

The reconstructions estimate the large-scale (low- ℓ) spherical harmonic coefficients, $a_{\ell m}$, by treating the signal at small scales as noise and only considering data external to the sky cut. If, as with the cosmic microwave background (CMB), the field to be reconstructed is not band-limited, the proliferation of small-scale sig-

nal makes the reconstruction noisy to the point of being useless. Input maps are therefore smoothed – necessarily prior to cutting the sky – to truncate the signal and remove sources of confusion below a chosen angular scale [9]. However, smoothing leaks contamination from the masked region into the trusted data [10, 11], and the reconstructed spherical harmonic coefficients, $\hat{a}_{\ell m}$, are biased. In this work we explore the causes and expected magnitudes of this bias, and discuss how it can be mitigated.

MAXIMUM-LIKELIHOOD RECONSTRUCTION

We begin with a description of the standard implementation of maximum-likelihood CMB $a_{\ell m}$ reconstruction. The first step of the reconstruction process is to band-limit the temperature field by smoothing, typically with a Gaussian kernel of width 10° FWHM. As this removes information on the smallest scales, the map resolution can be downgraded to reduce computation time. The $\hat{a}_{\ell m}$ s in the range $2 \leq \ell \leq \ell_{\text{max,rec}}$ (represented for ease as the $n_{a_{\ell m}}$ -element vector $\hat{\mathbf{a}}$, where $n_{a_{\ell m}} = (\ell_{\text{max,rec}} - 1)(\ell_{\text{max,rec}} + 3)$) are then reconstructed from the n_{pix} unmasked pixel temperatures, \mathbf{x} , using [3]

$$\hat{\mathbf{a}} = \mathbf{W}\mathbf{x}. \quad (1)$$

The reconstructed spherical harmonic coefficients maximize the likelihood of the residual noise, given the available data, if the reconstruction matrix, \mathbf{W} , is

$$\mathbf{W} = [\mathbf{Y}^t \mathbf{C}^{-1} \mathbf{Y}]^{-1} \mathbf{Y}^t \mathbf{C}^{-1}. \quad (2)$$

* stephen.feeney.09@ucl.ac.uk

† h.peiris@ucl.ac.uk

‡ apontzen@ast.cam.ac.uk

Here, \mathbf{Y} are the spherical harmonics calculated at each unmasked pixel¹, and \mathbf{C} is the pixel-space noise covariance matrix

$$C_{ij} = R_{ij} + \sum_{\ell=\ell_{\max, \text{rec}}+1}^{\ell_{\max}} \frac{2\ell+1}{4\pi} \bar{C}_\ell P_\ell(\hat{\mathbf{r}}_i \cdot \hat{\mathbf{r}}_j), \quad (3)$$

where \mathbf{R} is uncorrelated, low-amplitude regularizing noise added to prevent \mathbf{C} from becoming singular, \bar{C}_ℓ is the smoothed theory CMB angular power spectrum, and P_ℓ are the Legendre polynomials at unmasked pixels i, j . The sum over the multipoles $\ell_{\max, \text{rec}} < \ell \leq \ell_{\max}$ ensures that the small-angular-scale CMB multipoles we do not wish to reconstruct are treated as noise. As stated above the CMB power must be artificially truncated to restrict the number of ambiguous modes accessible to the reconstruction. The smoothing kernel is deconvolved from the $\hat{a}_{\ell m}$ s after reconstruction by dividing the $\hat{a}_{\ell m}$ s by the kernel's spherical harmonic transform.

If a foreground signal b_i is now introduced, so that $\mathbf{x} = \mathbf{Y}\mathbf{a} + \mathbf{b} + \mathbf{n}$ and \mathbf{a} is the CMB signal uncorrelated with \mathbf{b} and the noise \mathbf{n} , the mean and variance of the reconstruction error $\epsilon_{\ell m} = \hat{a}_{\ell m} - a_{\ell m}$ are

$$\langle \epsilon \rangle = \mathbf{W}\mathbf{b} \quad (4)$$

and

$$\langle \epsilon \epsilon^t \rangle - \langle \epsilon \rangle \langle \epsilon^t \rangle = \mathbf{W}\mathbf{C}\mathbf{W}^t, \quad (5)$$

respectively.

Throughout this work, we reconstruct the spherical harmonic coefficients up to $\ell_{\max, \text{rec}} = 10$. The noise covariance matrix includes CMB power in the range $\ell_{\max, \text{rec}} < \ell \leq \ell_{\max} = 32$ unless explicitly stated; this value is chosen such that modes with $\ell > \ell_{\max}$ are suppressed to $\mathcal{O}(\text{few } \%)$ by the smoothing. The WMAP 5-year best-fit C_ℓ s [12] are chosen for the theory CMB angular power spectrum². Input maps are smoothed at HEALPix [13] resolution $N_{\text{side}} = 512$ before being downgraded to $N_{\text{side}} = 16$ to retain the information required for the reconstruction while minimizing the number of pixels included in the noise covariance matrix. Diagonal regularizing noise \mathbf{R} is added at the level of $2 \mu\text{K}^2$ to allow the inversion of the noise covariance matrix despite the presence of some null modes (which are irrelevant to the reconstruction).

SMOOTHING-INDUCED BIAS

This section outlines how a bias arises from smoothing-induced contamination of the unmasked pixels [10, 11].

For clarity, we illustrate the smoothing-induced bias in the maximum-likelihood reconstruction (and, later, our proposed solutions) with results plotted in both harmonic and pixel space.

It is beyond the scope of this work to estimate accurate foreground residuals resulting from different component separation methods; instead, we choose some residuals for the purpose of illustration. Following Ref. [14], the residual foregrounds are taken to be 1% of the difference between the WMAP 7-year Internal Linear Combination (ILC) [15] and V-band [16] temperature maps. The resulting map is *indicative* of the extent and amplitude of the residual foregrounds in the WMAP 7-year ILC, the data-set for which the smoothing-induced bias was first described [9–11]. It is important to note that a considerably higher level of contamination is present in the foreground-reduced maps for each individual WMAP frequency band provided by the WMAP team: contaminants of ~ 50 times those used here are visible in these maps. The residual foregrounds are restricted to the pixels within the sky cut, which in this work masks only the Galaxy and not individual point-sources. These degree-scale point-source cuts are, unsurprisingly, not found to significantly bias the large-scale reconstructed $\hat{a}_{\ell m}$ s, and so for clarity Galaxy-only contaminants and masks are considered.

The addition of simulated residuals and the spherical harmonic transform are linear, so the smoothing-induced bias is given by the reconstructed $\hat{a}_{\ell m}$ s of the simulated residual map. We take the “standard” 10° -FWHM-Gaussian-smoothed maximum-likelihood reconstruction, using data outside the Galaxy-only part of the WMAP 7-year KQ85 mask, as our fiducial maximum-likelihood estimator (hereafter the “Gaussian ML” $\hat{a}_{\ell m}$ s). The Gaussian ML $\hat{a}_{\ell m}$ s generated from the simulated residual foregrounds are plotted (deep-blue solid line) in Fig. 1, along with the full-sky $a_{\ell m}$ s (dotted line) for comparison. The smoothing-induced leakage from within the sky cut is clear to see: the reconstruction picks up about half of the power in the foreground residuals, even though the simulated residuals are entirely confined to the sky cut.

The largest peaks in the bias affect the $\hat{a}_{\ell m}$ s satisfying $\ell = 2n, m = 0$ for integer n (in Fig. 1, HEALPix index $\ell^2 + \ell + m = \{6, 20, 42 \dots\}$) [3], and are positive for odd n , negative for even n . The pattern of these peaks can be explained by examining the reconstruction of the simulated Galactic residuals, plotted in Fig. 2, which are coldest along the Galactic plane. Smoothing these residuals reduces the pixel temperature values approximately symmetrically around the Galactic mask, and therefore pollutes the azimuthal modes which are also symmetric about the equator. The bias is positive for modes which have minima at the equator, and negative for those with maxima. There are also secondary peaks at $\ell = 2n+1, m = 1$ (in Fig. 1, $\ell^2 + \ell + m = \{13, 31, 57 \dots\}$), which again are positive for odd n , negative for even n .

¹ Without loss of generality, the reconstruction matrix in this work is formed from the *real* spherical harmonics.

² Our results are not sensitive to the small differences between different WMAP releases in the best-fit cosmology.




FIG. 1. The spherical harmonic coefficients of the simulated foreground residuals, calculated using the Gaussian ML reconstruction (deep-blue solid line), the full-sky data (dotted line) and the 10° Top-Hat ML reconstruction (light-blue solid line on x -axis). It is clear that the Gaussian ML reconstruction leaks around half of the information from within the sky cut; this can be counteracted by smoothing with a top-hat kernel and using an extended mask. The HEALPix index $\ell^2 + \ell + m$ maps each ℓ, m combination to a unique index into the array of $a_{\ell m}$ s.

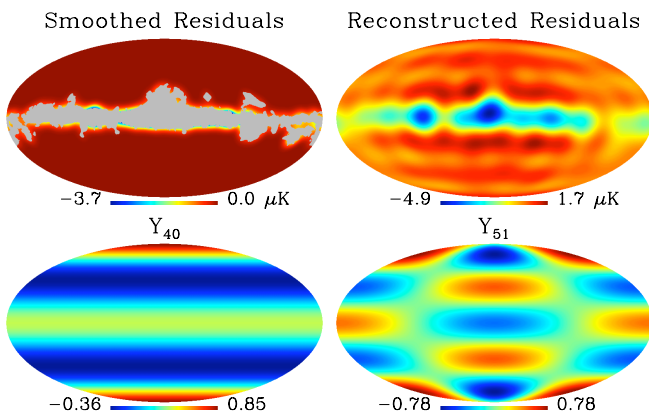


FIG. 2. Clockwise, from top left: simulated Galactic foregrounds, smoothed by a 10° -FWHM Gaussian and masked with the Galaxy-only KQ85 mask; the $\ell \leq 10$ Gaussian ML $\hat{a}_{\ell m}$ s reconstructed from the simulated Galactic foregrounds; the real spherical harmonics Y_{40} and Y_{51} . The simulated foregrounds yield a negative bias in the Y_{40} mode, and a positive bias in the Y_{51} mode.

These modes pick out the concentration of reconstructed foreground power in the Galactic centre.

The leakage of information from within the sky cut can also be demonstrated in pixel-space. Taking the WMAP 7-year ILC map, the full-sky $a_{\ell m}$ s are extracted, and the Gaussian ML reconstruction is performed. The spherical harmonic coefficients in each case are then used to reconstruct the input ILC map using only $2 \leq \ell \leq 10$, as plotted in Fig. 3. The maps formed from the full-

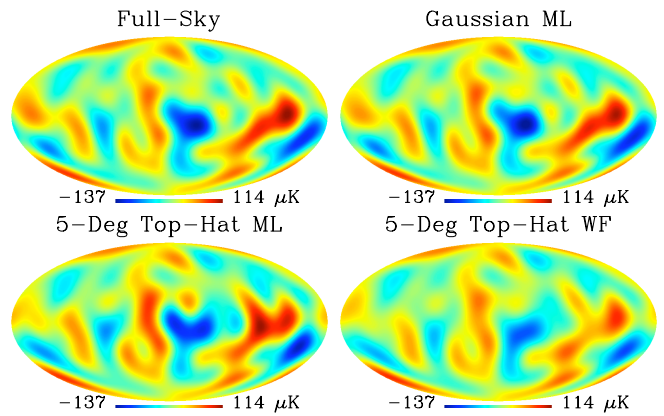


FIG. 3. The 7-year ILC $2 \leq \ell \leq 10$ modes, plotted using (clockwise from top-left) full-sky $a_{\ell m}$ s, Gaussian ML $\hat{a}_{\ell m}$ s, 5° Top-Hat ML $\hat{a}_{\ell m}$ s, and Top-Hat WF $\hat{a}_{\ell m}$ s.

sky- $a_{\ell m}$ s (top-left) and reconstructed from Gaussian-ML- $\hat{a}_{\ell m}$ s (top-right) are almost identical, even in the Galactic plane, confirming that the reconstruction has access to information well inside the sky cut.

Although the bias illustrated in Fig. 1 looks problematic, its amplitude is at least reduced over using the contaminated full sky, and so the simple Gaussian-ML procedure may yet turn out to be useful. To compare with other possible approaches, we first need to discuss its standard deviation – i.e. the scatter induced by the $\ell > \ell_{\text{max,rec}}$ modes – which is calculated using Eq. 5 and plotted (as the deep-blue narrowest band) in Fig. 4. This is a few μK at most. For a given ℓ , the modes that are reconstructed with the least precision are those with $|m| = \ell$, with the $m = \ell$ modes typically the worst. This confirms the observation in Pontzen and Peiris (2010) [14] that the sky cut removes the most information from modes with power concentrated towards the equator, and particularly those with extrema at $\phi = 0^\circ$, where the mask is at its widest [11]. The mask is plotted for reference in Fig. 5, along with examples of the affected modes. We see that, typically, three modes per ℓ will have increased bias or variance, but for most modes both the mean reconstruction error and its variance will be small.

In order to compare estimators, we must first quantify their performance over the range of multipoles considered. The performance measure

$$z_{\ell m} = \langle \epsilon_{\ell m}^2 \rangle = \text{mean}(\epsilon_{\ell m})^2 + \text{var}(\epsilon_{\ell m}) \quad (6)$$

provides the expected size of the reconstruction error $\epsilon_{\ell m}$ for each mode: summing over all modes

$$Z = \sum_{\ell, m} z_{\ell m} \quad (7)$$

therefore yields a complete measure of each estimator's performance. Any alternative estimator which removes




FIG. 4. The one-standard-deviation ranges of the reconstruction error $\epsilon_{\ell m}$ for the Gaussian ML $\hat{a}_{\ell m}$ s (narrowest, deep-blue band), 5° Top-Hat ML $\hat{a}_{\ell m}$ s (widest, light-blue band), and Top-Hat WF $\hat{a}_{\ell m}$ s (intermediate, mid-blue band).




FIG. 5. Top: the WMAP 7-year Galaxy-only KQ85 mask (light-blue central region) extended by 2.5° (mid-blue) and 5° (dark-blue). Bottom: the real spherical harmonics Y_{44} and Y_{4-4} . The concentration of $Y_{\ell\pm\ell}$ mode power towards the equator results in increased estimator variance in those modes.

the smoothing-induced bias should be preferred only if its Z value is lower than that of the Gaussian ML reconstruction, and, indeed, the contaminated full-sky $a_{\ell m}$ s. In fact, using the 1% ILC-V foreground residuals, $Z \simeq 33 \mu K^2$ using the full-sky $a_{\ell m}$ s, compared to $\sim 265 \mu K^2$ for the Gaussian ML estimator. If the residual foregrounds employed in this work are an accurate reflection of those present in the WMAP 7-year ILC map, then the contaminated full-sky $a_{\ell m}$ s provide a better estimate of the cosmological signal than the Gaussian ML reconstruction. The second form of Eq. 6 shows that both bias (see Fig. 1) and variance (see Fig. 4) in the recon-




FIG. 6. Top: the residual foreground map employed in this work (left), and a simple model comprising a 3° -wide $-20 \mu K$ equatorial band. Bottom: the full-sky $a_{\ell m}$ s (dotted) and reconstructed Gaussian ML $\hat{a}_{\ell m}$ s (solid) of the residual foreground maps. The simple band model captures most of the features of the smoothing-induced bias injected by the more-complex residuals.

struction increase the value of Z . The variance term is independent of the contamination, while the bias scales linearly with the contamination. Therefore, if the amplitude of residual foregrounds in the ILC map is higher than in our illustrative example, the reduction in bias due to the use of the Gaussian ML $\hat{a}_{\ell m}$ s will eventually overcome the variance introduced by the reconstruction. For residual levels 3 – 4 times higher than those used here, the Gaussian ML reconstruction should be used instead of the contaminated full-sky $a_{\ell m}$ s. However, as we have seen, the Gaussian ML estimator (as implemented thus far) does not eliminate the bias due to smoothing-related leakage of contaminants from within the masked region.

While our simulated foreground residuals are simply meant to be indicative, we nevertheless expect that the smoothing bias is mainly sensitive to the *amplitude* of the residuals, and not their precise morphology. This can be seen in Fig. 6, where we have modeled the residuals as a simple bar in the Galactic plane, while rescaling the amplitude to match our 1%(ILC-V) model. This highly simplified model is able to capture most of the features of the bias in harmonic space, as seen in the lower panels of this Figure.

ELIMINATING THE BIAS

At this stage, we are presented with something of a conundrum: smoothing is essential to the reconstruction process, but it is exactly this smoothing that is biasing the results. The simplest solution to this issue is to remove the areas of the sky that are within one smoothing scale of the main Galactic sky cut. However, the smooth-

ing kernel typically used in the standard reconstruction algorithm is a Gaussian, with support across the full sky in pixel space, and the set of contaminated pixels is hence poorly defined. This problem can be solved by using a kernel with finite pixel-space support – for simplicity we choose a top-hat – as all contaminated pixels fall within a kernel’s radius of the mask.

Naïvely selecting the diameter of the top-hat smoothing kernel to be 10° , we reconstruct the simulated foreground residuals using the Galaxy-only KQ85 mask extended by 5° – hereafter the “ 10° Top-Hat ML” reconstruction. The results are plotted in light blue (pale line along x -axis) in Fig. 1: the smoothing-induced bias has been eliminated. However, the measure of reconstruction quality has deteriorated dramatically to $Z \simeq 11\,252\,\mu\text{K}^2$, significantly worse than the Gaussian ML reconstruction. There are two reasons for this increase in Z , which is now sourced entirely by increased variance in the reconstruction. Firstly, the top-hat smoothing kernel has support over a greater range of multipoles than the Gaussian kernel, and so more ambiguous modes contribute to the covariance matrix (in this case, and for all further kernels, we increase ℓ_{max} to 1024 to capture all relevant modes, even though at 10° the smoothing kernel is effectively band-limited at $\ell \sim 100$).³ The second factor is that the reconstruction-error variance increases very rapidly with the area of the sky that is masked [3]. This suggests the use of narrower smoothing kernels, although this necessarily increases the power of the high- ℓ noise.

The interplay between the variance injected by decreased smoothing and increased masking is shown in Fig. 7. Here, the total reconstruction-error variance is plotted for top-hat smoothing kernels of diameter $3^\circ - 10^\circ$, and hence mask extensions of $1.5^\circ - 5^\circ$. The reduction in variance due to minimizing the sky cut dominates the added noise from narrower smoothing. We therefore select the width of the smoothing kernel to be as small as possible, given the resolution of the degraded map. At $N_{\text{side}} = 16$, the pixels are $\sim 3^\circ$ across, so to avoid injecting bias through pixelization [11] (which would not be captured by Eq. 5) we choose our optimal kernel diameter to be just larger: 5° . Hereafter, we refer to this reconstruction – using the 5° top-hat smoothing kernel and Galaxy-only KQ85 mask extended by 2.5° – as the “ 5° Top-Hat ML” reconstruction.

The reconstruction-error variance calculated for the 5° Top-Hat reconstruction is plotted as the light-blue outermost region of Fig. 4. Even using the minimum possible mask extension the reconstruction-error variance is still an order of magnitude larger than that of the Gaussian ML estimator. In terms of the measure of recon-




FIG. 7. The impact of mask extension and smoothing-kernel diameter on the quality of the bias-free maximum-likelihood reconstruction. Reconstruction-error variances are calculated using top-hat smoothing kernels of varying diameter, and extending the KQ85 mask by one kernel radius each time. Smoothing at the lowest-possible scale will produce the most faithful reconstruction. Note that the curve is not smooth as both the extent and *shape* of the mask change as it is extended.

struction quality, the biased Gaussian ML reconstruction ($Z \sim 265\,\mu\text{K}^2$) should be strongly preferred to the 5° Top-Hat ML case ($Z \sim 8466\,\mu\text{K}^2$) for residual levels comparable to those used in this work. As the Top-Hat ML reconstruction is unbiased, this value of Z is fixed (for a given sky cut). Thus, only if the residuals are greater than $\sim 25\%$ ILC–V will the 5° Top-Hat ML reconstruction outperform the Gaussian ML reconstruction. Note that the quality of the reconstruction could be improved further if it was performed at higher resolution, as smaller smoothing kernels could be used. This will necessarily have to be traded off against the increased computational requirements.

REDUCING THE VARIANCE

The increase in variance encountered when using extended sky cuts is far beyond that expected due to the reduction in pixel count: reducing f_{sky} from 81% to 74% should, assuming uncorrelated pixels for simplicity, increase the variance by only $\sim 10\%$. The dominant issue is that the maximum-likelihood reconstruction allows the temperature field in the masked region of the sky to have infinite variance. For small sky cuts (and small $\ell_{\text{max, rec}}$) this is fine: one cannot “hide” large-scale power within the cut, and so the variance on the large-scale $\hat{a}_{\ell m}$ s is low. Extending the KQ85 mask not only increases its overall width, but also closes a number of small gaps that allow the estimator limited access to the poorly-constrained equatorial modes (see Fig. 5). The estimator is therefore

³ Truncating the kernel by selecting a lower ℓ_{max} leads to $\mathcal{O}(10\%)$ variations in the results.

free to fill the cut with significant low- ℓ power (compare Fig. 3 top-right and bottom-left), and the estimator variance rises rapidly.

The variance of the reconstruction error can be reduced by enforcing a prior on the power within the Galactic cut using Wiener-filtering (Method 5 in Ref. [7]; see also Refs. [4–6, 8]). The Wiener-filtered $\hat{a}_{\ell m}$ s then maximize the posterior probability of reconstructing the underlying $a_{\ell m}$ s, given the trusted data and theoretical power spectrum. In practice, this is achieved by adding a theoretical (inverse) covariance matrix for the multipoles of interest to the reconstruction matrix

$$\mathbf{W} = [\mathbf{S}^{-1} + \mathbf{Y}^t \mathbf{C}^{-1} \mathbf{Y}]^{-1} \mathbf{Y}^t \mathbf{C}^{-1}. \quad (8)$$

\mathbf{S} here is a diagonal $n_{a_{\ell m}} \times n_{a_{\ell m}}$ matrix with elements equal to the smoothed theory power spectrum \bar{C}_ℓ . The theory prior restricts this “Top-Hat WF” reconstruction from filling the sky cut with arbitrary power (see Fig. 3 bottom-right), and reduces the variance of the estimator’s reconstruction error accordingly (see the mid-blue intermediate region in Fig. 4). The analytic expectation for the reconstruction-error variance is

$$\langle \epsilon \epsilon^t \rangle - \langle \epsilon \rangle \langle \epsilon^t \rangle = \mathbf{S} [\mathbf{S} + \mathbf{N}]^{-1} \mathbf{N}, \quad (9)$$

where \mathbf{N} is the noise power spectrum, defined to be $\mathbf{N} = [\mathbf{Y}^t \mathbf{C}^{-1} \mathbf{Y}]^{-1}$ (*i.e.* the variance of the Top-Hat ML reconstruction error). By adding a theory prior to the power within the sky cut, and hence requiring finite power in that region, the Wiener-filtered reconstruction tends to produce $\hat{a}_{\ell m}$ s that are closer to zero than the maximum-likelihood case. While this could be seen as biasing the $\hat{a}_{\ell m}$ s toward lower values⁴, it can also be interpreted as being conservative, and applying the prior belief that the information within the mask is similar to the trusted information outside the mask. In other words, we should be happy to trade off a small multiplicative bias against a significant reduction in variance.

This is automatically encapsulated in the measure of reconstruction quality Z for the 5° Top-Hat WF estimator, which has improved to $\sim 1521 \mu\text{K}^2$. However this is still worse than that of the Gaussian ML estimator. The bias arising in the Top-Hat WF reconstruction is not from smoothing but from a prior, so Z is fixed for a given mask, and always lower than that of the corresponding Top-Hat ML reconstruction. For contamination levels of $\gtrsim 10\%$ ILC–V (such as those found in the foreground-reduced maps for the individual WMAP frequency bands), the 5° Top-Hat WF reconstruction therefore represents the most reliable estimator considered in this work.

⁴ Assuming for clarity zero noise, the Wiener-filtered reconstruction yields $\hat{\mathbf{a}} = \mathbf{S}[\mathbf{S} + \mathbf{N}]^{-1} \mathbf{a}$, *i.e.* a multiplicative bias. Note that the ensemble average $\langle \hat{\mathbf{a}} \rangle = \langle \mathbf{a} \rangle = 0$.

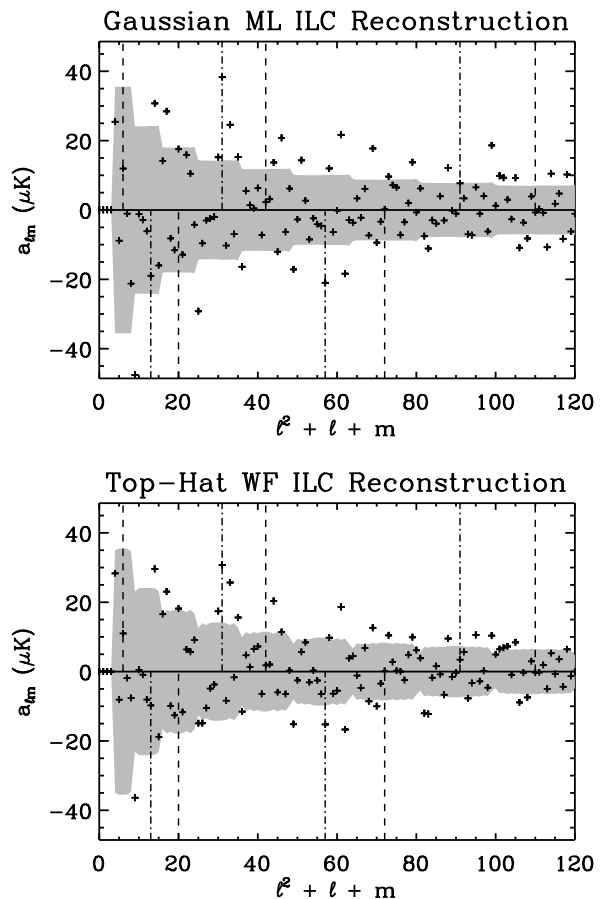


FIG. 8. The reconstructed WMAP 7-year ILC $\hat{a}_{\ell m}$ s, calculated using the Gaussian ML reconstruction (top) and the 5° Top-Hat WF reconstruction (bottom). The shaded areas represent the estimator standard deviations. The modes that are most contaminated by the simulated foregrounds in the Gaussian ML reconstruction are indicated, along with their expected sign, by dashed ($Y_{\text{even}0}$) and dash-dotted ($Y_{\text{odd}1}$) lines.

We do not know the precise magnitude or morphology of the residual foregrounds in the WMAP data. We can, however, examine the Gaussian ML and Top-Hat WF reconstructions of the ILC by eye to determine if there is an obvious impact due to residual foregrounds. The $\hat{a}_{\ell m}$ s generated from these two reconstructions are plotted in Fig. 8, along with the *estimator* standard deviation ($\sqrt{\langle |\hat{a}_{\ell m}|^2 \rangle - \langle \hat{a}_{\ell m} \rangle^2}$). Comparing the two plots, we see that there is very little difference between the $\hat{a}_{\ell m}$ s returned in each case. Further, no modes look statistically anomalous at the 3σ -level, even those that we expect to be contaminated from the simple residuals model used here.

RELATION OF $a_{\ell m}$ RECONSTRUCTION TO THE QML ESTIMATOR FOR THE C_ℓ S

We have so far discussed estimating the full-sky $a_{\ell m}$ s from cut-sky data, which is equivalent to reconstructing the smoothed temperature field. However, in the context under which the smoothing-induced bias was revealed [9–11] it is in fact only the angular power spectrum C_ℓ of the temperature field which is required.

A popular method for estimating the full-sky angular power spectrum is to adopt the quadratic maximum-likelihood estimator as first derived in Ref. [17]. It has been noted (see e.g. Section 3 of Ref. [9] for a complete discussion) that the QML estimator can be formed using the maximum-likelihood $\hat{a}_{\ell m}$ s. On the surface, the QML estimates (henceforth denoted \hat{C}_ℓ) may therefore seem to be susceptible to similarly problematic contamination from a smoothing stage.

However, this is not the case: in fact the \hat{C}_ℓ s are far more robust to the content of the cut because the smoothing can be conducted on vastly smaller scales (e.g. 1° in Pontzen and Peiris (2010) [14]). Note that Copi *et al.* (2011) [11] miss this point, because they consider only two extreme cases: (i) smoothing at 10° and (ii) failing to smooth. They therefore reach the erroneous conclusion that the QML estimator is susceptible to contamination from within the mask. We explicitly verified that the pipeline used by Pontzen and Peiris (2010) [14] is independent of any contamination placed fully inside the mask.

The above paragraphs at first appear to be contradictory, since they simultaneously claim (a) that the QML power spectrum estimates can be formed out of the ML temperature field reconstruction; and (b) that the QML power spectrum estimates can still be constructed from maps smoothed on degree scales (whereas the $\hat{a}_{\ell m}$ s will necessarily become noisy for sufficiently high $\ell_{\text{max,rec}}$). However, this is not a true contradiction because the QML estimates are not formed directly from the noisy $\hat{a}_{\ell m}$ s, but rather through an expression (Eq. 23 of Ref. [9]) which specifically downweights poorly constrained modes. It is this cautious treatment of ambiguous modes which makes power spectrum estimation, as opposed to $a_{\ell m}$ reconstruction, so well-behaved, irrespective of the shape of the smoothing kernel employed.

DISCUSSION

Maximum-likelihood estimators, $\hat{a}_{\ell m}$, are often used to reconstruct the large-scale spherical harmonic coefficients, $a_{\ell m}$, from partial-sky data. The technique relies on smoothing to restrict the amount of small-scale noise accessible to the reconstruction, but smoothing has been shown to contaminate “clean” pixels with residual foregrounds from within the sky cut. In this work, we have

examined the impact of this smoothing-induced bias on the maximum-likelihood reconstruction. We have shown that it is possible to mitigate the bias by removing the contaminated regions, but these are only well-defined if smoothing is performed using a kernel with finite support on the sky. This precludes the use of the commonly used Gaussian kernel. Cutting a larger portion of the sky greatly increases the variance of the reconstruction, but it is possible to counteract this effect by enforcing a prior on the reconstructed coefficients using a Wiener filter. We have therefore proposed an estimator – using top-hat smoothing, extended masks and a Wiener-filtered reconstruction – which does not suffer from smoothing-induced bias. By considering the expectation of the square of the reconstruction error, $Z = \sum_{\ell, m} \langle (\hat{a}_{\ell m} - a_{\ell m})^2 \rangle$, we have compared the performance of the maximum-likelihood and Wiener-filtered estimators in the presence of simulated CMB foreground residuals.

The reconstruction performance measure Z scales with the estimators’ bias and variance, which in turn are governed by the amplitude of contamination and the size of the sky cut, respectively. The fiducial maximum-likelihood reconstruction is performed using relatively small sky cuts, but is susceptible to contamination through smoothing-induced bias; the finite-smoothing Wiener-filtered reconstruction does not suffer from smoothing-induced bias, but makes use of extended masks. Increasing the level of contamination therefore increases Z for the maximum-likelihood reconstruction only, which suggests that there is a level of contamination above which one should switch from the maximum-likelihood to the Wiener-filtered reconstruction.

Given an estimate of the morphology and amplitude of the contaminants within the cut sky, one can predict which modes will be biased and by how much, and hence determine the threshold at which one should swap estimators. We find that this threshold is relatively insensitive to the precise morphology of foreground residuals at large scales, and is mainly governed by their amplitude. Calculating Z for the two estimators in the presence of estimated foreground residuals, we determine this threshold to be ~ 10 times the amplitude of the foreground residuals used in this work. Assuming that the ILC contains similar levels of contamination to those used here, we therefore recommend the use of either the contaminated full-sky $a_{\ell m}$ s or the fiducial maximum-likelihood $\hat{a}_{\ell m}$ s when handling this data-set. However, when using foreground-reduced maps for individual WMAP frequencies, which contain much greater foreground residuals, the Wiener-filtered reconstruction will provide the best estimate of the large-scale underlying CMB signal. Note that, as the Wiener-filtered $\hat{a}_{\ell m}$ s are a maximum-posterior solution, care must be taken if the reconstruction output is being used for further model-selection steps. The reconstruction techniques are, however, most commonly used to test the null hypothesis, in which case

the prior employed in this work is completely appropriate.

For problems requiring only a power spectrum (as opposed to the full temperature field) the issues described in this paper are essentially irrelevant because the smoothing can be conducted on vastly smaller scales, the resulting range of poorly constrained modes being automatically downweighted.

We thank George Efstathiou for interesting discussions. SMF is supported by the Perren Fund and STFC. HVP is supported by Marie Curie grant MIRG-CT-2007-203314 from the European Commission, and by STFC and the Leverhulme Trust. AP is supported by Emmanuel College, Cambridge. We acknowledge use of the HEALPix package and the Legacy Archive for Microwave Background Data Analysis (LAMBDA). Support for LAMBDA is provided by the NASA Office of Space Science.

[1] R. K. Sachs and A. M. Wolfe, *Astrophys. J.* **147**, 73 (1967).
 [2] A. Slosar, C. Hirata, U. Seljak, S. Ho, and N. Padmanabhan, *JCAP* **0808**, 031 (2008), arXiv:0805.3580 [astro-ph].
 [3] A. de Oliveira-Costa and M. Tegmark, *Phys. Rev.* **D74**, 023005 (2006), arXiv:astro-ph/0603369.

[4] N. Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (The MIT Press, 1964).
 [5] E. F. Bunn *et al.*, *Astrophys. J.* **432**, L75 (1994), arXiv:astro-ph/9404007.
 [6] S. Zaroubi, Y. Hoffman, K. B. Fisher, and O. Lahav, *Astrophys. J.* **449**, 446 (1995), arXiv:astro-ph/9410080.
 [7] M. Tegmark, *Astrophys. J.* **480**, L87 (1997), arXiv:astro-ph/9611130.
 [8] P. Bielewicz, K. M. Gorski, and A. J. Banday, *Mon. Not. Roy. Astron. Soc.* **355**, 1283 (2004), arXiv:astro-ph/0405007.
 [9] G. Efstathiou, Y.-Z. Ma, and D. Hanson, (2009), arXiv:0911.5399 [astro-ph.CO].
 [10] R. Aurich and S. Lustig, *Mon. Not. Roy. Astron. Soc.* **411**, 124 (2011), arXiv:1005.5069 [astro-ph.CO].
 [11] C. J. Copi, D. Huterer, D. J. Schwarz, and G. D. Starkman, (2011), arXiv:1103.3505 [astro-ph.CO].
 [12] M. R. Nolta *et al.* (WMAP), *Astrophys. J. Suppl.* **180**, 296 (2009), arXiv:0803.0593 [astro-ph].
 [13] K. M. Gorski *et al.*, *Astrophys. J.* **622**, 759 (2005), arXiv:astro-ph/0409513.
 [14] A. Pontzen and H. V. Peiris, *Phys.Rev.D* **81**:103008, 2010 (2010), 1004.2706.
 [15] B. Gold *et al.*, *Astrophys. J. Suppl.* **192**, 15 (2011), arXiv:1001.4555 [astro-ph.GA].
 [16] N. Jarosik *et al.*, *Astrophys. J. Suppl.* **192**, 14 (2011), arXiv:1001.4744 [astro-ph.CO].
 [17] M. Tegmark, *Phys. Rev.* **D55**, 5895 (1997), arXiv:astro-ph/9611174.